Machine Characterizations of the Classes of the W-Hierarchy

Yijia Chen and Jörg Flum

Abteilung für Mathematische Logik, Universität Freiburg, Eckerstr. 1,
79104 Freiburg, Germany.
`chen@zermelo.mathematik.uni-freiburg.de` `Joerg.Flum@math.uni-freiburg.de`

Abstract. We give machine characterizations of the complexity classes of the W-hierarchy. Moreover, for every class of this hierarchy, we present a parameterized halting problem complete for this class.

1 Introduction

Parameterized complexity theory provides a framework for a refined complexity analysis of algorithmic problems that are intractable in general. Central to the theory is the notion of *fixed-parameter tractability*, which relaxes the classical notion of tractability, polynomial time computability, by admitting algorithms whose runtime is exponential, but only in terms of some *parameter* that is usually expected to be small. As a complexity theoretic counterpart, a theory of *parameterized intractability* has been developed. In classical complexity theory, the notion of NP-completeness is central to a nice and simple theory for intractable problems. Unfortunately, the world of parameterized intractability is more complex: there is a big variety of seemingly different classes of parameterized intractability. Nevertheless, it can be argued that the classes $W[1], W[2], \ldots$ of the W-hierarchy together with some other classes like $W[P]$ correspond to NP in classical complexity theory. In particular, all these classes are defined by parameterized variants of the NP-complete satisfiability problem. Unfortunately, the definition of these classes by means of complete problems makes it not easy to understand them. The authors of [4] tried to remedy this situation by presenting machine characterizations for some of the classes. But as they remarked "it remains an interesting open problem to find natural machine characterizations for the classes $W[t]$ for $t \geq 2$". In this paper we obtain such characterizations.

By definition, a parameterized problem is fixed-parameter tractable, if it is decidable by a (deterministic) algorithm in at most $f(k) \cdot p(n)$ steps for some computable function f and some polynomial p. Here k denotes the size of the parameter and n the size of the input. Problems in any of the "intractable" parameterized complexity classes mentioned so far, also are decidable in at most $f(k) \cdot p(n)$ steps but by a *nondeterministic* algorithm. Thus, a first attempt to characterize one of these intractable classes by machines consists in considering nondeterministic algorithms that perform at most $f(k) \cdot p(n)$ steps but restricting further the number of nondeterministic steps. This approach led to the following

results in [4]: A problem is in W[P] if and only if it is decidable in at most $f(k) \cdot p(n)$ steps by an algorithm whose number of nondeterministic steps is bounded in terms of the parameter. Similarly, a problem is in W[1], if in addition the nondeterministic steps are performed at the end of the computation.

Already in [4], nondeterministic random access machines turned out to be the appropriate machine model in order to get clear formulations of the characterizations. In their nondeterministic steps these machines are able to guess natural numbers $(\leq f(k) \cdot p(n))$; these numbers are considered as names of objects. To obtain machine characterizations of the classes of the W-hierarchy we have to consider the corresponding alternating random access machines, but at the same time we have to ensure that the programs only have access to (the properties of) the elements named by the guessed numbers and not to the numbers themselves.

In [3], Cesati and Di Ianni prove that the halting problem p-HPNMT for nondeterministic multitape Turing machines, parameterized by the number of steps, is W[2]-hard by reducing the parameterized dominating set problem to it: the machine first guesses the elements of a dominating set and then checks that they really constitute a dominating set. An analysis of the use of and the access to the guessed elements in this algorithm helped the authors to find the machine characterizations of the classes of the W-hierarchy. In Section 4 we present a (short) proof of membership in W[2] of p-HPNMT by reducing it to a model-checking problem in W[2], a result obtained in [2] by different means.

As the corresponding proof in [4] shows, the machine characterization of W[1] is closely related to the W[1]-completeness of the halting problem for non-deterministic Turing machines. For $t \geq 2$, the W[t]-complete halting problem we present in the last section refers to alternating Turing machines with oracles and it is not so closely related to the corresponding machine characterization but to a logical model-checking problem complete for W[t].

2 Preliminaries

In this section we recall some definitions and results and fix our notations.

2.1. Relational Structures and First-Order Logic. A *vocabulary* τ is a finite set of relation symbols. Each relation symbol has an *arity*. A *(relational) structure* $\mathcal{A}$ of vocabulary τ, or τ-structure, consists of a set A called the *universe*, and an interpretation $R^{\mathcal{A}} \subseteq A^r$ of each r-ary relation symbol $R \in \tau$. We synonymously write $\bar{a} \in R^{\mathcal{A}}$ or $R^{\mathcal{A}}\bar{a}$ to denote that the tuple $\bar{a} \in A^r$ belongs to the relation $R^{\mathcal{A}}$. For example, we view a *directed graph* as a structure $\mathcal{G} = (G, E^{\mathcal{G}})$, whose vocabulary consists of one binary relation symbol E. $\mathcal{G}$ is an (undirected) *graph*, if $E^{\mathcal{G}}$ is irreflexive and symmetric.

The class of all first-order formulas is denoted by FO. They are built up from atomic formulas using the usual boolean connectives and existential and universal quantification. Recall that *atomic formulas* are formulas of the form $x = y$ or $Rx_1 \ldots x_r$, where $x, y, x_1, \ldots, x_r$ are variables and R is an r-ary relation

symbol. For $t \geq 1$, Σ_t denotes the class of all first-order formulas of the form

$$\exists x_{11} \ldots \exists x_{1k_1} \forall x_{21} \ldots \forall x_{2k_2} \ \ldots \ Q x_{t1} \ldots Q x_{tk_t} \ \psi,$$

where $Q = \forall$ if t is even and $Q = \exists$ otherwise, and where ψ is quantifier-free. Π_t-formulas are defined analogously starting with a block of universal quantifiers. Let $t, u \geq 1$. A formula φ is $\Sigma_{t,u}$, if it is Σ_t and all quantifier blocks after the leading existential block have length $\leq u$.

If Φ is a class of formulas, then $\Phi[\tau]$ denotes the class of all formulas of vocabulary τ in Φ. If $\mathcal{A}$ is a τ-structure and $\varphi \in \mathrm{FO}[\tau]$ a sentence, i.e., a formula without free variables, then we write $\mathcal{A} \models \varphi$ to denote that $\mathcal{A}$ satisfies φ.

The proof of the following lemma is easy.

Lemma 1. *Let $\varphi_1(\bar{x}), \ldots, \varphi_m(\bar{x})$ and $\psi_1(\bar{x}, \bar{y}), \ldots, \psi_m(\bar{x}, \bar{y})$ be formulas in $\mathrm{FO}[\tau]$, where $\bar{x} = x_1 \ldots x_r$ and $\bar{y} = y_1 \ldots y_s$ are sequences of variables that have no variable in common. Assume $Q_1, \ldots, Q_r, Q'_1, \ldots, Q'_s \in \{\forall, \exists\}$. If $\mathcal{A}$ is a τ-structure with $\mathcal{A} \models \forall \bar{x} \neg(\varphi_i \wedge \varphi_j)$ for $i \neq j$, then $\mathcal{A}$ satisfies both or none of the formulas*

$$Q_1 x_1 \ldots Q_r x_r \bigwedge_{i=1}^{m} (\varphi_i \to Q'_1 y_1 \ldots Q'_s y_s \psi_i),$$

$$Q_1 x_1 \ldots Q_r x_r Q'_1 y_1 \ldots Q'_s y_s \bigwedge_{i=1}^{m} (\varphi_i \to \psi_i).$$

2.2. Parameterized Complexity. A *parameterized problem* is a set $Q \subseteq \Sigma^* \times \Pi^*$, where Σ and Π are finite alphabets. If $(x, y) \in \Sigma^* \times \Pi^*$ is an instance of a parameterized problem, we refer to x as the *input* and to y as the *parameter*. We usually denote the length of the input string x by n and the length of the parameter string y by k.

For example, for a class S of structures and a class L of first-order formulas,

$$\text{p-MC(S, L)} := \{(\mathcal{A}, \varphi) \mid \mathcal{A} \in \mathrm{S}, \ \varphi \text{ a sentence in L, and } \ \mathcal{A} \models \varphi\}$$

is the *parameterized model-checking problem for* S *and* L; mostly, for easier readability, we present parameterized problems in the following form:

> p-MC(S,L) *Input:* $\mathcal{A} \in \mathrm{S}$.
> *Parameter:* φ, a sentence in L.
> *Problem:* $\mathcal{A} \models \varphi$?

Definition 1. *A parameterized problem $Q \subseteq \Sigma^* \times \Pi^*$ is* fixed-parameter tractable, *if there are a computable function $f : \mathbb{N} \to \mathbb{N}$, a polynomial p, and an algorithm that, given a pair $(x, y) \in \Sigma^* \times \Pi^*$, decides if $(x, y) \in Q$ in at most $f(k) \cdot p(n)$ steps.*

FPT *denotes the complexity class consisting of all fixed-parameter tractable parameterized problems.*

Complementing the notion of fixed-parameter tractability, there is a theory of parameterized intractability. It is based on the following notion of reduction:

Definition 2. *An* FPT-reduction *from the parameterized problem $Q \subseteq \Sigma^* \times \Pi^*$ to the parameterized problem $Q' \subseteq (\Sigma')^* \times (\Pi')^*$ is a mapping $R : \Sigma^* \times \Pi^* \to (\Sigma')^* \times (\Pi')^*$ such that:*

1. *For all $(x, y) \in \Sigma^* \times \Pi^*$: $(x, y) \in Q \iff R(x, y) \in Q'$.*
2. *There exists a computable function $g : \mathbb{N} \to \mathbb{N}$ such that for all $(x, y) \in \Sigma^* \times \Pi^*$, say with $R(x, y) = (x', y')$, we have $k' \leq g(k)$ (where $k = |y|$ and $k' = |y'|$).*
3. *There exist a computable function $f : \mathbb{N} \to \mathbb{N}$ and a polynomial p such that R is computable in time $f(k) \cdot p(n)$.*

We write $Q \leq^{\mathrm{FPT}} Q'$ if there is an FPT-reduction from Q to Q' and set $[Q]^{\mathrm{FPT}} := \{Q' \mid Q' \leq^{\mathrm{FPT}} Q\}$. For a class C of parameterized problems, we let $[\mathrm{C}]^{\mathrm{FPT}} := \bigcup_{Q \in \mathrm{C}} [Q]^{\mathrm{FPT}}$.

Denote by GRAPH the class of all finite graphs and by STR the class of all finite structures. The parameterized complexity classes $\mathrm{W}[1], \mathrm{W}[2], \dots$ of the W-hierarchy are defined as the closure of a family of parameterized problems under FPT-reductions. For $\mathrm{W}[t]$, the defining family of problems consists of parameterized versions of the satisfiability problems for circuits of *weft t* (and varying depth).

For the purposes of this paper, the most appropriate way to introduce the complexity classes of the W-hierarchy is the following (cf. [7], [8]):

Definition 3. *For $t \geq 1$, $\mathrm{W}[t]$ is the class of all parameterized problems that, for some $u \geq 1$, are* FPT-*reducible to* p-MC(GRAPH, $\Sigma_{t,u}$), *that is,*

$$\mathrm{W}[t] = [\{\text{p-MC}(\text{GRAPH}, \Sigma_{t,u}) \mid u \geq 1\}]^{\mathrm{FPT}}.$$

The following equivalent characterization of $\mathrm{W}[t]$ is well-known:

Proposition 1. $\mathrm{W}[t] = [\{\text{p-MC}(\text{STR}, \Sigma_{t,u}[\tau]) \mid u \geq 1,\ \tau \text{ vocabulary}\}]^{\mathrm{FPT}}$.

At various places of this paper we tacitly make use of the following remark:

Remark 1. Sometimes, in order to show that a given parameterized problem is in $\mathrm{W}[t]$, we will present a reduction to p-MC(STR, $\Sigma_{t,u}[\tau]$) for some vocabulary τ also containing a fixed finite number of constant symbols or even, we will consider a reduction where the number of constants depends on the parameter. This will allow to express properties in a more readable fashion. If we have a fixed finite number of constant symbols, they can be eliminated by using appropriate unary relations that are singletons; then, in the $\Sigma_{t,u}$-formula to be defined, the constants are replaced by variables that are existentially quantified (in the first block) and get their right value using the corresponding relations. If the number of constants symbols depends on the parameter, e.g., we use constants for $0, 1, \dots, k$, in order to stay within a fixed vocabulary (independent of k) these constants can be eliminated by a singleton relation for 0 and the binary successor relation on $\{0, 1, \dots, k\}$ and again by existentially quantified variables.

Sometimes we refer to the complexity classes of the A-hierarchy (cf. [8]):

Definition 4. *For $t \geq 1$, $\mathrm{A}[t] = [\text{p-MC}(\text{GRAPH}, \Sigma_t)]^{\mathrm{FPT}}$.*

3 Machine Characterization

In [4], machine characterizations of the classes W[1], A[t] for $t \geq 1$, and W[P] using nondeterministic and alternating random access machines (RAMs) were presented. The nondeterministic RAMs are based on the standard random access machines (cf. [10]). The model was non-standard when it came to nondeterminism. Instead of allowing the machines to nondeterministically choose one bit, or an instruction of the program to be executed next, the authors allowed them to nondeterministically choose a natural number, more precisely, there was an additional instruction "GUESS i j" whose semantics was: Guess a natural number less than or equal to the number stored in register i and store it in register j. In [4] it was remarked: "While this form of nondeterminism may seem unnatural at first sight, we would like to argue that it is very natural in many typical 'applications' of nondeterminism. For example, a nondeterministic algorithm for finding a clique in a graph guesses a sequence of vertices of the graph and then verifies that these vertices indeed form a clique. Such an algorithm is much easier described on a machine that can guess the numbers representing the vertices of a graph at once, rather than guessing their bits."

The alternating RAMs, in addition to the instructions of the form "GUESS i j" (denoted by "EXISTS i j" in the context of alternating machines) also had "FORALL i j" instructions. The semantics was defined as usually for alternating machines. The computations of alternating machines suitable for A[t] have a parameter-bounded final part which contains all the EXISTS- and FORALL-instructions and have at most $t-1$ alternations (see [4] for the precise statement). For W[t] it seems not to suffice (see Section 6) to bound the length of the blocks without alternation, but we have to restrict the access to the numbers guessed in the EXISTS- and FORALL-instructions: as just remarked in the algorithm for finding a clique, we view these numbers as labels of certain objects and the type of machine we are going to introduce only has access to the properties of these objects and not directly to the labels.

We turn to the precise definition of the random access machines we are going to use and that we call W-RAMs. A W-RAM has the

- the *standard registers* $0, 1, \ldots$, their contents are denoted by $r_0, r_1, \ldots$, respectively.
- the *guess registers* $0, 1, \ldots$, their contents are denoted by $g_0, g_1, \ldots$, respectively.

All registers have initial value 0. Often we denote g_{r_i}, i.e., the contents of the guess register whose index is the content of the ith standard register, by $g(r_i)$.

The W-RAM has all the standard instructions for the standard registers (cf. Section 2.6 of [10]; e.g., the arithmetic operations are addition, subtraction, and

division by two (rounded off)). Moreover, it has four additional instructions:

Instruction	Semantics
EXISTS $\uparrow j$	guess a natural number $\leq r_0$; store it in the r_jth guess register
FORALL $\uparrow j$	guess a natural number $\leq r_0$; store it in the r_jth guess register
JG= $i\ j\ c$	if $g(r_i) = g(r_j)$, then jump to the instruction with label c
JG0 $i\ j\ c$	if $r_{\langle g(r_i), g(r_j)\rangle} = 0$, then jump to the instruction with label c.

Here, $\langle\ ,\ \rangle : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ is any simple coding of ordered pairs of natural numbers by natural numbers such that $\langle i, j\rangle \leq (1 + \max\{i, j\})^2$ and $\langle 0, 0\rangle = 0$. Of course, the semantics of the EXISTS- and FORALL-instructions are the same, but we view them as existential and universal instructions, respectively. All other instructions are said to be deterministic. If the machine stops, it *accepts* its input, if $r_0 = 0$, otherwise it rejects it.

The following lemma, whose proof is immediate, is crucial for the main theorem of this section; it shows that the contents of the standard registers depend only on the sequence of executed instructions:

Lemma 2. *Assume that, for a given input, we have two (partial) computations on a W-RAM. If the same sequence of instructions is carried out in both computations, then the contents of the standard registers will be the same.*

Definition 5. *A program $\mathbb{P}$ for a W-RAM is an* AW-program, *if there is a computable function f and a polynomial p such that for every input (x, y) with $|x| = n$ and $|y| = k$ the program $\mathbb{P}$ on every run*

1. *performs at most $f(k) \cdot p(n)$ steps;*
2. *at most $f(k)$ steps are existential or universal;*
3. *at most the first $f(k) \cdot p(n)$ standard registers are used;*
4. *at every point of the computation the registers contain numbers $\leq f(k) \cdot p(n)$.*

The promised machine characterization of W[t] reads as follows:

Theorem 1. *Let Q be a parameterized problem and $t \geq 1$. Then Q is in W[t] if and only if there is a $u \geq 1$ and there are a computable function h and an AW-program $\mathbb{P}$ for a W-RAM such that $\mathbb{P}$ decides Q and such that for every run of $\mathbb{P}$ on an instance (x, y) of Q as input (with $|y| = k$)*

- *all existential and universal steps are among the last $h(k)$ steps of the computation,*
- *there are at most $t - 1$ alternations between existential and universal states, and the first guess step is existential,*
- *every block without alternations, besides the first one, contains at most u guess steps.*

6

Proof. Assume first that $Q \in W[t]$, then $Q \leq^{\mathrm{FPT}}$ p-MC(GRAPH, $\Sigma_{t,u}$) for some $u \geq 1$. Hence there are computable functions f and g, a polynomial $p \in \mathbb{N}[x]$, and an algorithm $\mathbb{A}$ assigning to every (x, y), in time $\leq f(k) \cdot p(n)$, a graph $\mathcal{G} = \mathcal{G}_{x,y}$ and a sentence $\varphi = \varphi_{x,y} \in \Sigma_{t,u}$, say,

$$\varphi = \exists x_{11} \ldots \exists x_{1k_1} \forall x_{21} \ldots \forall x_{2k_2} \ldots Q x_{t1} \ldots Q x_{tk_t} \psi,$$

with $k_2, \ldots, k_t \leq u$, with $|\varphi| \leq g(k)$, and with a quantifier-free ψ, such that

$$Qxy \iff \mathcal{G} \models \varphi.$$

The claimed AW-program $\mathbb{P}$ for a W-RAM, on input (x, y), proceeds as follows:

1. It computes the graph $\mathcal{G} = (G, E^{\mathcal{G}})$, say with $G = \{1, \ldots, m\}$, and stores its adjacency matrix in the standard registers: $r_{\langle i,j \rangle} = 0 \iff E^{\mathcal{G}} ij$.
2. It computes φ.
3. It checks whether $\mathcal{G} \models \varphi$.

To carry out point 3, the program $\mathbb{P}$, using the EXISTS- and FORALL-instructions guesses the values of the quantified variables. Then, it checks the quantifier-free part using the JG=- and JG0-instructions. The number of steps needed for point 3 can be bounded by $h(k)$ for some computable h. Hence, all existential and universal steps are among the last $h(k)$ steps of the computation.

   Now assume that $\mathbb{P} = (\pi_1, \ldots, \pi_m)$ is an AW-program deciding Q and that $u \geq 1$, h, and $\mathbb{P}$ have the properties stated at the right side of the equivalence claimed in the theorem. For the program $\mathbb{P}$ choose the function f and the polynomial p according to Definition 5. We claim that $Q \in W[t]$. By Proposition 1, it suffices to show that $Q \leq^{\mathrm{FPT}}$ p-MC(STR, $\Sigma_{t,2\cdot u}[\tau]$) for some τ. The set of instruction numbers of $\mathbb{P}$ is $\{1, \ldots, m\}$, more precisely, π_i is the instruction of $\mathbb{P}$ with instruction number i. We denote instruction numbers (= potential contents of the program counter) by $c, c_1, \ldots$ and finite sequences of instruction numbers by $\bar{c}$. $\ell(\bar{c})$ denotes the last instruction number of the sequence $\bar{c}$, and $[\bar{c}]$ the sequence obtained from $\bar{c}$ by omitting its last member. Fix an instance (x, y) of Q. Let

$$C := \bigcup_{0 \leq r \leq h(k)-1} \{1, \ldots, m\}^r \quad \text{and} \quad N := \{0, 1, \ldots, f(k) \cdot p(n)\}.$$

with $k = |y|$. We look for a structure $\mathcal{A}$ and a $\Sigma_{t,2\cdot u}$-sentence φ such that

$$\mathbb{P} \text{ accepts } (x, y) \iff \mathcal{A} \models \varphi.$$

Let $\bar{c}_0$ be the sequence of instruction numbers of the deterministic part of $\mathbb{P}$ on input (x, y) ending with the instruction number of the first existential instruction. As universe A of $\mathcal{A}$ we take $A := C \cup N$. Moreover, in $\mathcal{A}$ there are the binary relation $\leq^{\mathcal{A}}$, the natural ordering on N, and ternary relations $R^{\mathcal{A}}$ and $T^{\mathcal{A}}$ defined by

$$R^{\mathcal{A}} \bar{c} i j \iff \bar{c} \in C, \ i, j \in N, \text{and if } \mathbb{P}, \text{ on input } (x, y), \text{ carries out}$$

$$\text{the sequence of instructions } [\bar{c}_0\,\bar{c}), \text{ then } r_i = j.$$

$$T^{\mathcal{A}}\bar{c}ij \iff \bar{c} \in C, \ i, j, \langle i, j \rangle \in N, \text{and if } \mathbb{P}, \text{ on input } (x, y), \text{ carries out}$$
$$\text{the sequence of instructions } [\bar{c}_0\,\bar{c}), \text{ then } r_{\langle i,j \rangle} = 0.$$

Moreover, we have a constant for 0. This finishes the definition of $\mathcal{A}$, which can be constructed within the time allowed by an FPT-reduction.

We turn to the definition of φ. First, we fix $\bar{c} \in C$: Let $i = i(\bar{c})$ be the number of blocks without alternations of the sequence of instructions determined by $[\bar{c}_0\bar{c})$; let $j = j(\bar{c})$ be the number of guesses in the ith block. If $i \le t$ and if each block, besides the first one, has length $\le u$, we introduce a formula

$$\varphi_{\bar{c}}(\bar{x}_1, \ldots, \bar{x}_{i-1}, x_{i,1} \ldots x_{i,j}) \tag{1}$$

where $\bar{x}_1 := x_{1,1}, \ldots, x_{1,h(k)}$ and $\bar{x}_s := x_{s,1} \ldots x_{s,u}$ for $s = 2, \ldots, i-1$ with the intuitive meaning

> if a partial run of $\mathbb{P}$ has $\bar{c}_0\,\bar{c}$ as sequence of instructions numbers and if every variable $x_{i',j'}$ displayed in (1) has, as value, the j'th guess of the i'th block (and $x_{i',j'} = 0$, if there was no such guess), then there is an accepting continuation of this run.

Then, for the empty sequence $\emptyset$ of instruction numbers and $\varphi := \varphi_\emptyset$, we have

$$\mathbb{P} \text{ accepts } (x, y) \iff \mathcal{A} \models \varphi.$$

For $\bar{c} \in C$ of maximal length, $|\bar{c}| = h(k) - 1$, we set (recall that, by definition, a computation accepts its input, if $r_0 = 0$ at the end of the computation)

$$\varphi_{\bar{c}} := \begin{cases} \text{TRUE,} & \text{if } \pi_{\ell(\bar{c})} = \text{STOP (i.e., } \pi_{\ell(\bar{c})} \text{ is the STOP-instruction) and } R\bar{c}\,0\,0 \\ \text{FALSE,} & \text{otherwise.} \end{cases}$$

If $\bar{c} \in C$ and $|\bar{c}| < h(k) - 1$, we assume that $\varphi_{\bar{c}'}$ has already been defined for all $\bar{c}'$ with $|\bar{c}| < |\bar{c}'|$. The definition depends on the type of the instruction of $\mathbb{P}$ with instruction number $\ell(\bar{c})$.

If $\pi_{\ell(\bar{c})} = \text{STOP}$, then again

$$\varphi_{\bar{c}} := \begin{cases} \text{TRUE,} & \text{if } R\bar{c}\,0\,0 \\ \text{FALSE,} & \text{otherwise.} \end{cases}$$

The definition of $\varphi_{\bar{c}}$ is simple for the standard instructions, e.g., if $\pi_{\ell(\bar{c})} = \text{STORE} \uparrow u$ (i.e., "$\pi_{\ell(\bar{c})} = r_{r_u} := r_0$"), then $\varphi_{\bar{c}} := \varphi_{\bar{c}\,\ell(\bar{c})+1}$.

We give the definitions for the new instructions: If $\pi_{\ell(\bar{c})} = \text{EXISTS} \uparrow v$, then (for $i = i(\bar{c})$ and $j = j(\bar{c})$)

$$\varphi_{\bar{c}} := \begin{cases} \exists x_{i,j+1} \exists y (R\bar{c}\,0\,y \land x_{i,j+1} \le y \land \varphi_{\bar{c}\,\ell(\bar{c})+1}), & \text{if } i = 1 \text{ or } (i \text{ is odd and } j < u) \\ \exists x_{i+1,1} \exists y (R\bar{c}\,0\,y \land x_{i+1,1} \le y \land \varphi_{\bar{c}\,\ell(\bar{c})+1} \\ \qquad \land x_{i,j+1} = 0 \land \ldots \land x_{i,u} = 0), & \text{if } i \text{ is even, } i < t, \text{ and } j \le u \\ \text{FALSE,} & \text{otherwise.} \end{cases}$$

The definition is similar for instructions of the type FORALL $\uparrow v$, but then the variables are quantified universally.

Assume $\pi_{\ell(\bar{c})} = \text{JG} = v\ w\ c$. We need $g(r_v)$ and $g(r_w)$. Determine the actual contents v_0 and w_0 of the vth and the wth standard register, i.e., v_0 and w_0 with $R^{\mathcal{A}}\bar{c}vv_0$ and $R^{\mathcal{A}}\bar{c}ww_0$. Consider the sequence of instructions given by $\bar{c}$ and determine the last instructions in it of the form FORALL $\uparrow z$ or EXISTS $\uparrow z$ such that at that time $r_z = v_0$, say, it is the j_0th guess in the i_0th block. Similarly, let the j_1th guess in the i_1th block be the last instruction of the form FORALL $\uparrow z$ or EXISTS $\uparrow z$ such that at that time $r_z = w_0$ (the case that such instructions do no exist is treated in the obvious way). Then set

$$\varphi_{\bar{c}} := (x_{i_0,j_0} = x_{i_1,j_1} \to \varphi_{\bar{c}\,c}) \wedge (\neg x_{i_0,j_0} = x_{i_1,j_1} \to \varphi_{\bar{c}\,\ell(\bar{c})+1}).$$

Assume $\pi_{\ell(\bar{c})} = \text{JG0}\ u\ v\ c$. As in the preceding case, let x_{i_0,j_0} and x_{i_1,j_1} denote the actual values of the r_uth and the r_vth guess register, respectively. Then set

$$\varphi_{\bar{c}} := (T\bar{c}x_{i_0,j_0}x_{i_1,j_1} \to \varphi_{\bar{c}\,c}) \wedge (\neg T\bar{c}x_{i_0,j_0}x_{i_1,j_1} \to \varphi_{\bar{c}\,\ell(\bar{c})+1}).$$

As already mentioned above, we set $\varphi := \varphi_\emptyset$. By Lemma 1, one easily verifies that φ is equivalent to a $\Sigma_{t,2\cdot u}$-formula. Clearly, the size of φ can be bounded in terms of $h(k)$ and

$$Qxy \iff \mathbb{P} \text{ accepts } (x,y)$$
$$\iff \mathcal{A} \models \varphi,$$

which gives the desired reduction showing that $Q \in \text{W}[t]$. $\square$

4   W[2] and multitape machines

Among the many known W[1]-complete problems, of course, the most generic one is the halting problem p-HPN for nondeterministic Turing machines (cf. [1]):

---
p-HPN   *Input:* A nondeterministic Turing machine M.
    Parameter: $k \in \mathbb{N}$.
        Problem: Does M accept the empty word in at most k steps?

Surprisingly, the same problem for nondeterministic Turing machines with several tapes is W[2]-complete (cf. [3] and [2]). Mike Fellows pointed out this result to the second author and encouraged him to look for a machine characterization of W[2]. In this section we present a simple proof that the halting problem for multitape machines is in W[2]; it avoids the equality $\text{W}[2] = \text{W}^*[2]$ (cf. [5]), used in [2] and for which no simple proof is known.

Proposition 2. *For some vocabulary τ, p-HPNMT $\leq^{\text{FPT}}$ p-MC(STR, $\Sigma_{2,3}[\tau]$).*

Here p-HPNMT denotes the halting problem for nondeterministic multitape Turing machines:

Proof. Let M be a nondeterministic Turing machine with w_0 (work) tapes. We aim at a structure $\mathcal{A} = \mathcal{A}_{M,k}$ and a $\Sigma_{2,3}$-sentence $\varphi = \varphi_{M,k}$ such that

$$(M, k) \in \text{p-HPNMT} \iff \mathcal{A} \models \varphi.$$

Let Σ and Q be the alphabet and the set of states of M, respectively. The instructions of M have the form

$$q\,(a_1, \ldots, a_{w_0}) \to q'(a'_1, \ldots, a'_{w_0})\,(h_1, \ldots, h_{w_0})$$

where $q, q' \in Q$, $a_1, \ldots, a_{w_0}, a'_1, \ldots, a'_{w_0} \in \Sigma \cup \{*\}$ ($*$ is the blank symbol) and $h_1, \ldots, h_{w_0} \in \{-1, 0, 1\}$. Let T be the set of tuples $(b_1, \ldots, b_{w_0}) \in (\Sigma \cup \{*\})^{w_0}$ and H the set of tuples $(h_1, \ldots, h_{w_0}) \in \{-1, 0, 1\}^{w_0}$ occurring in instructions.

The structure $\mathcal{A}$ has the universe

$$A := Q \cup (\Sigma \cup \{*\}) \cup \{-1, 0, 1, \ldots, \max\{w_0, k\}\} \cup T \cup H.$$

We need the natural ordering relation $\leq^{\mathcal{A}}$ on $\{-1, 0, 1, \ldots, \max\{w_0, k\}\}$, the 5-ary relation $D^{\mathcal{A}}$ (the "transition relation") and the ternary relation $P^{\mathcal{A}}$ (the "projection relation") defined by

$$D^{\mathcal{A}}qtq't'h \iff qt \to q't'h \text{ is an instruction of } M$$
$$P^{\mathcal{A}}wba \iff 1 \leq w \leq w_0,\ b \in T \cup H,\ b = (b_1, \ldots, b_{w_0}),\ \text{and } b_w = a.$$

Moreover, we have constant symbols for the initial state q_0, the accepting state q_{acc}, for $*$, and for $-1, w_0, 0, 1, \ldots, k$ (cf. Remark 1).

The formula φ we aim at, will express that there is an accepting run of length $\leq k$ (w.l.o.g. of length $= k$); among others, it will contain the variables $q_i, t_i, q'_i, t'_i, h_i$ for $i = 1, \ldots, k$, in fact, $q_i t_i \to q'_i t'_i h_i$ represents the ith instruction applied in the run; φ is obtained by existentially quantifying all variables in

$$\left(\varphi_{\text{init}}(q_1, t_1) \wedge \bigwedge_{i=1}^{k} D q_i\, t_i\, q'_i\, t'_i\, h_i \wedge \bigwedge_{i=1}^{k-1} q'_i = q_{i+1} \wedge q_k = q_{\text{acc}} \wedge \psi \right),$$

where $\varphi_{\text{init}}(q_1, t_1) := (q_1 = q_0 \wedge \forall w(1 \leq w \leq w_0 \to Pwt_1*))$ and where ψ is a universal formula expressing that the sequence of instructions can be applied: For this purpose, for $i = 1, \ldots, k$, we introduce quantifier-free formulas

$$\varphi_i^L(w, p, x, \bar{v}_i) \text{ and } \varphi_i^S(w, p, \bar{v}_i)$$

with $\bar{v}_i := q_1, t_1, q'_1, t'_1, h_1, \ldots, q_{i-1}, t_{i-1}, q'_{i-1}, t'_{i-1}, h_{i-1}$ and with the meaning

10

if, starting with the empty tape, the sequence of instructions $\bar{v}_i$ has been carried out, then the pth cell of the wth tape contains the letter x,

and

if, starting with the empty tape, the sequence of instructions $\bar{v}_i$ has been carried out, then the head of the wth tape scans the pth cell,

respectively. Then, as ψ, we can take

$$\psi := \forall w \forall p \forall x \bigwedge_{i=1}^{k} \Big((\varphi_i^L(w,p,x,\bar{v}_i) \wedge \varphi_i^S(w,p,\bar{v}_i)) \to Pwt_i x \Big).$$

The simultaneous definition of φ_i^L and φ_i^P by induction on i is routine, e.g.,

$$\varphi_1^L(w,p,x) := (1 \le w \le w_0 \wedge 1 \le p \le k \wedge x = *);$$

$$\varphi_{i+1}^L(w,p,x,\bar{v}_{i+1}) := (1 \le w \le w_0 \wedge 1 \le p \le k) \wedge$$
$$((\neg \varphi_i^S(w,p,\bar{v}_i) \wedge \varphi_i^L(w,p,x,\bar{v}_i)) \vee (\varphi_i^S(w,p,\bar{v}_i) \wedge Pwt_i'x)).$$

One easily verifies that φ is (logically equivalent to) a $\Sigma_{2,3}$-sentence. □

Corollary 1. *The halting problem* p-HPNMT *for multitape machines is* W[2]-*complete.*

Proof. By Proposition 1 and the preceding proposition, p-HPNMT is in W[2]. To show that p-HPNMT is W[2]-hard we recall the proof of [3] that the parameterized dominating set problem p-DS can be reduced to p-HPNMT. The essential problem is to obtain a corresponding machine in the time allowed by an FPT-reduction.

Suppose $(\mathcal{G}, k)$ is an instance of p-DS with $G = \{a_1, \ldots, a_n\}$. Let the multitape machine M have $n+1$ tapes, numbered by 0 to n, and let $\Sigma := G \cup \{\text{yes}\}$ be its alphabet. In the first step all heads move one cell to the right. In the next $2 \cdot k$ steps, only the 0th head is active: it (nondeterministically) writes k elements of G on its tape, say $b_1, \ldots, b_k$ (the elements of the intended dominating set), and goes back to the cell containing b_1. In the next k steps the 0th head again reads these elements; at the same time, in the jth step, the ith head checks whether $a_i = b_j$ or $Ea_i b_j$; in the affirmative case the ith head prints "yes" and moves to the right, in the negative case it neither prints nor moves; finally, the machine moves all heads one cell to the left and accepts, if the heads on the tapes number 1 to n read "yes". Clearly,

$$(\mathcal{G}, k) \in \text{p-DS} \iff (M, 3 \cdot (k+1)) \in \text{p-HPNMT}.$$ □

5 Complete halting problems for the W-hierarchy

As already mentioned at the beginning of Section 4, the halting problem for nondeterministic Turing machines, parameterized by the number of steps, is a

quite generic W[1]-complete problem. The corresponding halting problems for alternating machines yield complete problems for the classes of the A-hierarchy. Indeed, for $t \geq 1$, we have $A[t] = [\text{p-HPA}_t]^{\text{FPT}}$ (cf. [8]), where

---

p-HPA$_t$ *Input:* An alternating Turing machine M whose initial state is existential.

Parameter: $k \in \mathbb{N}$.

Problem: Does M accept the empty word in at most k steps with at most $t - 1$ alternations?

One would expect that in order to obtain complete problems for the classes of the W-hierarchy one has to bound the number of steps of all non-alternating blocks but the first one. More precisely, for $t, u \geq 1$, consider the following problem:

---

p-HPA$_{t,u}$ *Input:* An alternating Turing machine M whose initial state is existential.

Parameter: $k \in \mathbb{N}$.

Problem: Does M accept the empty word in at most k steps with at most $t - 1$ alternations, where every block of steps without alternation, besides the first one, has length $\leq u$?

Essentially along the lines of the proof of p-HPA$_t \in A[t]$ in [8], one can show that p-HPA$_{t,u} \in W[t]$. But the corresponding hardness proof does not seem to go through. Recall that in order to establish that p-MC(GRAPH, Σ_t) $\leq^{\text{FPT}}$ p-HPA$_t$, given a graph and a Σ_t-sentence φ, one constructs an alternating Turing machine that first associates values to the quantifiers in φ (the values for existentially quantified variables are chosen in existential states, the values for universally quantified variables in universal states) and that then checks whether the selected variables satisfy the quantifier-free part of φ, the *quantifier-free check*. Of course, a $\Sigma_{t,u}$-prefix would yield an alternation sequence according to p-HPA$_{t,u}$, but the number of steps needed for the quantifier-free check cannot be bounded in terms of t and u as required in p-HPA$_{t,u}$, it depends on φ. Therefore, we add suitable oracles to the Turing machines that carry out the quantifier-free check in a single step. Of course, we have to add them in such a way that the corresponding halting problem still is in W[t].

   As in an alternating machine, the set Q of states of an *alternating Turing machine M with oracle* is the disjoint union of the set of universal states Q_u and the set of existential states Q_e. The "oracle states" $q_?$, q_y, and q_n are all contained in Q_u or all in Q_e. Let Σ be the vocabulary of M and let $O \subseteq \Sigma^*$ be a language, the "oracle language". M^O, the machine M with oracle O, in state $q_?$ will check if the word to the left of the cell scanned by its head is in O and will change to the "yes state" q_y or to the "no state" q_n (without printing a letter nor moving its head).

For a graph $\mathcal{G} = (G, E^{\mathcal{G}})$ and a first-order formula $\psi(x_1, \ldots, x_p)$, let $O(\mathcal{G}, \psi)$ be the following set of words over G, $O(\mathcal{G}, \psi) \subseteq G^*$:

$$O(\mathcal{G}, \psi) := \{a_1 \ldots a_p \mid \mathcal{G} \models \psi(a_1, \ldots, a_p)\}.$$

Due to space limitations we omit the proof of the following theorem:

Theorem 2. *For $t \geq 1$,*

$$\mathrm{W}[t] = [\{\text{p-HPAO}_{t,u} \mid u \geq 1\}]^{\mathrm{FPT}},$$

where p-HPAO$_{t,u}$ is the parameterized halting problem for alternating Turing machines with oracle:

p-HPAO$_{t,u}$ *Input:*	An alternating Turing machine M with oracle whose initial state is existential and a graph $\mathcal{G} = (G, E^{\mathcal{G}})$ such that G is a subset of the alphabet of M.
Parameter:	$k \in \mathbb{N}$ and a quantifier-free formula $\psi(x_1, \ldots, x_p)$.
Problem:	Does $M^{O(\mathcal{G}, \psi)}$ accept the empty word in at most k steps with at most $t-1$ alternations, where every block of steps without alternation, besides the first one, has length $\leq u$?

6   Conclusions

Most standard complexity classes like LOGSPACE, NLOGSPACE, PTIME, NPTIME, the classes of the polynomial hierarchy, or PSPACE have definitions in terms of machines. By contrast, originally nearly all parameterized complexity classes containing intractable problems were defined via complete problems. In [4], machine characterizations of W[1], W[P], and of A[t] for $t \geq 1$ were presented; in this paper we derive such characterizations for W[t] for $t \geq 2$.

As mentioned at the beginning of Section 3, in [4] AW-programs for alternating RAMs were introduced, which have "unrestricted access" to the guessed numbers. For $t \geq 1$, denote by L[t] the class that satisfies Theorem 1 if we replace W-RAM by alternating RAM, i.e., L[t] is the class of parameterized problems Q such that there is a computable function h and an AW-program $\mathbb{P}$ for an alternating RAM deciding Q such that for every run of $\mathbb{P}$ on an instance (x, y) of Q as input (with $k = |y|$)

- all existential and universal steps are among the last $h(k)$ steps of the computation,
- there are at most $t-1$ alternations between existential and universal states, and the first guess step is existential,
- every block without alternations, besides the first one, contains at most u guess steps.

Clearly, by [4] and Theorem 1, we have

$$\mathrm{W}[t] \subseteq \mathrm{L}[t] \subseteq \mathrm{A}[t].$$

We do not know, if $\mathrm{W}[t] = \mathrm{L}[t]$ or if $\mathrm{L}[t] = \mathrm{A}[t]$. Let an *f-vocabulary* be a finite set of relation symbols, function symbols, and constant symbols. By an appropriate refinement of the proof of Theorem 1 one can show:

Theorem 3. *For $t \geq 1$,*

$$\mathrm{L}[t] = [\{\mathrm{p\text{-}MC}(\mathrm{STR}, \Sigma_{t,u}[\tau]) \mid u \geq 1, \tau \ f\text{-}vocabulary\}]^{\mathrm{FPT}}.$$

For every f-vocabulary τ, the classical problem $\mathrm{MC}(\mathrm{STR}, \Sigma_{t,u}[\tau])$ is in NP. Therefore, $\mathrm{FPT} \subseteq \mathrm{W}[t] \subseteq \mathrm{L}[t] \subseteq \text{para-NP}$ (compare [9] for the definition of para-NP). As $\mathrm{FPT} \neq \text{para-NP}$ is equivalent to $\mathrm{P} \neq \mathrm{NP}$ (cf. [9]), we obtain from the preceding theorem:

Corollary 2. *If $\mathrm{W}[t] \neq \mathrm{L}[t]$ for some $t \geq 1$, then $\mathrm{P} \neq \mathrm{NP}$.*

References

1. L. Cai, J. Chen, R.G. Downey, and M.R. Fellows. On the parameterized complexity of short computation and factorization. *Archive for Mathematical Logic*, 36:321–337, 1997.
2. M. Cesati. The Turing way to parameterized intractability, 2001. Submitted for publication.
3. M. Cesati and M. Di Ianni. Computation models for parameterized complexity. *Math. Log. Quart.*, 43:179–202, 1997.
4. Y. Chen, J. Flum, and M. Grohe. Bounded nondeterminism and alternation in parameterized complexity theory. To appear in *Proc. of the Conference on Computational Complexity*, 2003. Available at `http://www.dcs.ed.ac.uk/home/grohe/pub.html`.
5. R.G. Downey and M.R. Fellows. Threshold dominating set and an improved characterization of W[2]. *Theoretical Computer Science*, 209:123–140, 1996.
6. R.G. Downey and M.R. Fellows. *Parameterized Complexity*. Springer-Verlag, 1999.
7. R.G. Downey, M.R. Fellows, and K. Regan. Descriptive complexity and the *W*-hierarchy. In P. Beame and S. Buss, editors, *Proof Complexity and Feasible Arithmetic*, volume 39 of *AMS-DIMACS Volume Series*, pages 119–134. AMS, 1998.
8. J. Flum and M. Grohe. Fixed-parameter tractability, definability, and model checking. *SIAM Journal on Computing*, 31(1):113–145, 2001.
9. J. Flum and M. Grohe. Describing parameterized complexity classes. In H. Alt and A. Ferreira, editors, *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science*, volume 2285 of *Lecture Notes in Computer Science*, pages 359–371. Springer-Verlag, 2002.
10. C.H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.